

Zero-Shot Fine-Grained Image Classification Using Large Vision-Language Models

Md. Atabuzzaman Andrew Zhang Chris Thomas

Department of Computer Science
Virginia Tech

{atabuzzaman, azhang42, christhomas}@vt.edu

Abstract

Large Vision-Language Models (LVLMs) have demonstrated impressive performance on vision-language reasoning tasks. However, their potential for zero-shot fine-grained image classification, a challenging task requiring precise differentiation between visually similar categories, remains underexplored. We present a novel method that transforms zero-shot fine-grained image classification into a visual question-answering framework, leveraging LVLMs’ comprehensive understanding capabilities rather than relying on direct class name generation. We enhance model performance through a novel attention intervention technique. We also address a key limitation in existing datasets by developing more comprehensive and precise class description benchmarks. We validate the effectiveness of our method through extensive experimentation across multiple fine-grained image classification benchmarks. Our proposed method consistently outperforms the current state-of-the-art (SOTA) approach, demonstrating both the effectiveness of our method and the broader potential of LVLMs for zero-shot fine-grained classification tasks. Code and Datasets: <https://github.com/Atabuzzaman/Fine-grained-classification>

1 Introduction

Large Vision-Language Models (LVLMs) such as LLaVA (Liu et al., 2024a,b), BLIP-2 (Li et al., 2023), InstructBLIP (Dai et al., 2023), Molmo (Deitke et al., 2024), Qwen2-VL (Wang et al., 2024b), InternVL (Chen et al., 2024), and DeepSeek’s Janus-Pro (Chen et al., 2025) and so on, alongside closed-source models like GPT-4V (OpenAI, 2024) and others have revolutionized multimodal understanding. These models demonstrate remarkable zero-shot capabilities in tasks that require extensive real-world knowledge, such as Visual Question Answering (VQA), visual reasoning,

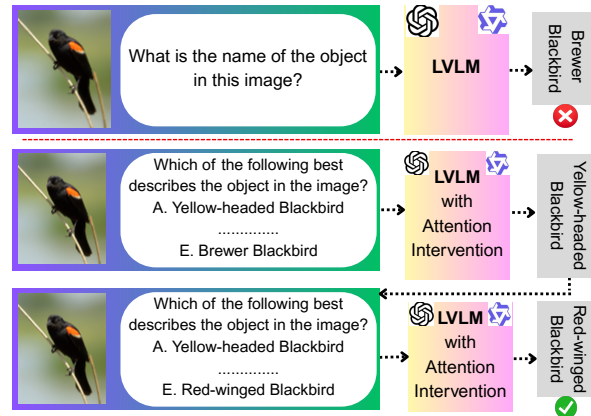


Figure 1: Overview of our zero-shot fine-grained image classification framework. Unlike existing approaches (top), which directly prompt a Large Vision-Language Model (LVLM) to generate a class name given an input image, our method leverages an LVLM combined with a proposed iterative multiple-choice question-answering strategy and an attention intervention technique to select the most accurate fine-grained class description. This framework effectively matches each input image with the most appropriate class description without requiring any training samples.

and image captioning. Their success largely stems from instruction-tuning and pre-training on vast multimodal datasets, enabling effective integration of visual content with textual knowledge.

However, despite their impressive performance in general vision-language tasks, LVLMs face significant challenges in fine-grained image classification, which requires distinguishing subtle visual differences between highly similar categories, such as bird species or car models (Kim and Ji, 2024). Unlike coarse-grained tasks where broad visual features suffice, fine-grained classification demands a precise understanding of nuanced visual characteristics and their alignment with detailed textual descriptions (Figure 1 & 2). This challenge becomes particularly acute in zero-shot settings, where models must identify fine-grained categories without task-specific training.

Initial attempts to address this challenge include *Finer* (Kim and Ji, 2024), which trained models to generate fine-grained attributes before predicting class names, and *GRAIN* (Halbe et al., 2024), which adapted CLIP (Radford et al., 2021) with LVLMs to measure similarity between generated text and ground truth class names. While these methods represent important first steps, their limited performance reveals critical gaps in current LVLM architectures and methodological approaches. Although LVLMs excel at generating descriptive captions and answering general questions about images, they struggle to capture and differentiate the subtle visual distinctions that define fine-grained categories (Kim and Ji, 2024). This limitation stems from their reliance on direct class name generation without sufficient domain knowledge and the use of less powerful models. These findings suggest that vision-language generation approaches may not adequately address the level of detail required for fine-grained classification tasks.

To address these limitations, we propose a novel method for zero-shot fine-grained image classification using LVLMs. Inspired by *VQAScore* (Lin et al., 2024), we first develop a strong foundation through likelihood-based yes/no QA classification. We then introduce an iterative multiple-choice question-answering (MCQA) approach to enable fine-grained classification. To reduce the computational overhead of iterative MCQA, we also explore an all-at-once approach that leverages the extended context capabilities of modern LVLMs to process all class descriptions simultaneously in a single forward pass. Rather than relying on direct class name generation, both approaches leverage the models’ comprehensive visual understanding capabilities. Our method aligns with traditional CLIP-based fine-grained classification approaches (Menon and Vondrick, 2022; Saha et al., 2024) by utilizing class labels and descriptions. Additionally, inspired by recent advances in LVLM hallucination reduction (Liu et al., 2025; Zhou et al., 2025; Huang et al., 2024; Wang et al., 2024a; Jiang et al., 2025), we introduce a simple yet effective attention intervention technique. This intervention integrates insights from multiple studies, combining early-layer visual information flow (Zhang et al., 2024) with deep-layer ground truth token detection capabilities (Wang et al., 2024a; Jiang et al., 2025) to enhance classification accuracy.

Furthermore, we identify a critical limitation in existing class descriptions, which are typically

generated using Large Language Models (LLMs) such as GPT-3/4 (Menon and Vondrick, 2022; Saha et al., 2024) or sourced from the internet and refined using GPT-4V (Kim and Ji, 2024). These descriptions often contain visual attributes that are absent from the actual class images or include irrelevant information that may impair model performance. To address this challenge, we develop curated class description benchmarks that more accurately capture the visual characteristics present in the images. Our approach investigates settings both with and without explicit class names in the descriptions, providing greater flexibility and robustness in classification tasks. Through extensive experiments on major fine-grained classification benchmark datasets, our methods demonstrate significant performance improvements over existing approaches and achieve SOTA performance. The main contributions of our work are as follows:

- We introduce a novel method for zero-shot fine-grained image classification that leverages LVLMs’ understanding capabilities rather than relying on class name generation.
- We propose an efficient attention intervention technique that enhances visual information flow by combining early-layer visual information with deep-layer semantic understanding.
- We introduce improved class description benchmarks that better align with visual features, addressing gaps in existing descriptions.
- We demonstrate substantial performance improvements over existing methods across five major fine-grained classification benchmarks.

2 Related Work

2.1 Zero-shot Image Classification using Vision-Language Models

Vision-Language Models (VLMs) are trained on massive web-scraped datasets for zero-shot image classification. CLIP (Radford et al., 2021) pioneered this approach through contrastive training of image and text encoders, optimizing for high similarity between matching pairs while minimizing similarity for non-matching ones. During inference, CLIP computes similarities between unseen images and text captions to perform classification. ALIGN (Jia et al., 2021) extended this framework by adding a data refinement stage.

Recent approaches leverage LLMs to generate richer class descriptions. Menon and Vondrick (2022) prompted LLMs to generate discriminative

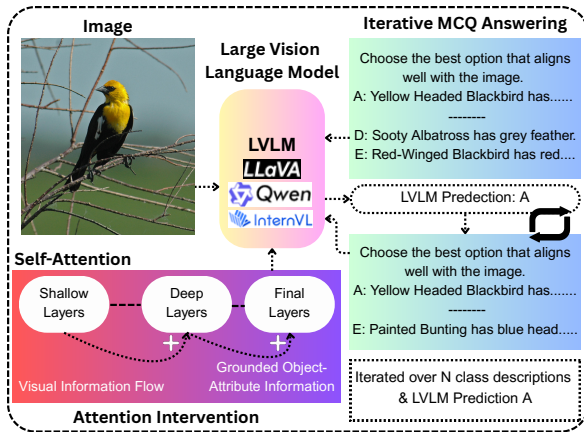


Figure 2: Overview of our proposed zero-shot fine-grained image classification framework using LVLm. The system takes an input image and class descriptions with a prompt, and uses an LVLm enhanced with an attention intervention mechanism. The framework employs an iterative MCQA approach where the LVLm selects the most appropriate class description through multiple rounds of refinement. The attention intervention module guides the visual information flow from shallow to deep layers, while deep layers provide grounded object-attribute information to final layers to improve classification accuracy.

attributes for each class, while Pratt et al. (2023) proposed generating multiple prompts per class to enrich text understanding. Saha et al. (2024) fine-tuned CLIP using attribute-based descriptions incorporating visual features, geographic information, and habitat context. Chiquier et al. (2024) enhanced interpretability by combining concept bottlenecks (Koh et al., 2020) with evolutionary search for LLM-generated concepts. These methods typically determine the predicted class through maximum cosine similarity between CLIP-generated text and image embeddings.

2.2 Zero-shot Image Classification using LVLms

Recent advancements in LVLms including LLaVA (Liu et al., 2024a,b), BLIP-2 (Li et al., 2023), InstructBLIP (Dai et al., 2023), and Qwen2-VL (Wang et al., 2024b) have depicted remarkable zero-shot capabilities in tasks like Visual Question Answering (VQA), visual reasoning, and image captioning. However, their ability to classify fine-grained images remains underexplored.

Prior work in zero-shot fine-grained classification with LVLms includes Finer (Kim and Ji, 2024), which generates class names and considers a prediction correct if the ground truth appears among the top 20 generated tokens. GRAIN (Halbe et al.,

2024) extended this approach by computing cosine similarity between CLIP embeddings of LLaVA-v1.6 generated text (Liu et al., 2024b) and ground truth class descriptions. In contrast, following traditional zero-shot classification paradigms, our work introduces a novel method that leverages class descriptions instead of generating class names. We provide these descriptions along with the image as input to LVLms. We either measure the probability of a ‘Yes’ response based on image-text alignment or prompt the model to select the most image-aligned description.

2.3 Attention Intervention

Recent research has explored various techniques to mitigate LVLm hallucinations in training-free settings. These approaches include visual contrastive decoding (VCD) (Leng et al., 2024), token selections (Huo et al., 2025), overconfident token penalization using attention weights (Huang et al., 2024; Gong et al., 2024), and attention intervention (Liu et al., 2025; Zhou et al., 2025). PAI (Liu et al., 2025) demonstrated that LVLms tend to underweight image tokens, with language decoders showing bias toward text tokens. To address this, PAI increased image token attention while subtracting language bias, following VCD (Leng et al., 2024). CAUSALMM (Zhou et al., 2025) employed various attention intervention strategies (random, uniform, reversed, and shift) and counterfactual reasoning at both visual and language attention levels to mitigate modality priors and enhance input-output alignment. DeCo (Wang et al., 2024a) proposed dynamic selection of previous layers to incorporate knowledge into final layers, while Jiang et al. (2025) introduced a system for hallucination mitigation through interpretation and editing of image embeddings. However, these methods are computationally expensive (Wang et al., 2024a; Jiang et al., 2025), as they rely on visual contrastive decoding and token or layer selection. This inefficiency makes these methods unsuitable for our classification task, which involves comparing a single image against multiple class descriptions. Instead, our method achieves greater efficiency by using a single forward pass while propagating visual information from shallow to deep layers and visual object-attribute grounded information (via attention weights) from deep layers to final layers. Our approach significantly reduces computational overhead while maintaining effectiveness in zero-shot fine-grained image classification.

3 Proposed Method

This section presents our proposed method for zero-shot fine-grained image classification using LVLMs. We begin by introducing a likelihood-based Yes/No Question-Answering (QA) approach, then propose an iterative MCQA framework to address the limitations of the initial approach. We introduce a simple yet effective attention intervention technique inspired by hallucination mitigation methods to enhance iterative MCQA approach performance. Figure 2 provides an overview of our iterative MCQA method with attention intervention. Finally, we discuss an all-at-once approach that leverages the extended context capabilities of modern LVLMs to reduce computational overhead.

3.1 Yes/No Question-Answering

Inspired by VQAScore (Lin et al., 2024), we introduce a likelihood-based Yes/No QA approach for zero-shot fine-grained image classification to determine the aligned class description for a given image. Given an image I , we formulate a query $Q(D_c)$ that prompts the LVLm to generate a binary response ("Yes" or "No") based on the alignment between a candidate class description D_c and the image. We use the probability of generating a "Yes" response as a confidence score to choose the best-matched class description for that image. Formally, for each candidate class $c \in \mathcal{C}$, we compute:

$$P_{\text{Yes}}(c) = p(\text{"Yes"} \mid I, Q(D_c)) \quad (1)$$

where $P_{\text{Yes}}(c)$ represents the probability of generating "Yes" for class c , indicating the model's confidence that D_c correctly describes the image.

The final predicted class \hat{y} is determined by selecting the class with the highest probability of generating "Yes":

$$\hat{y} = \arg \max_{c \in \mathcal{C}} P_{\text{Yes}}(c) \quad (2)$$

While this approach effectively selects class descriptions, it requires access to output token probabilities to compute the likelihood of generating "Yes", limiting its implementation to LVLms that provide token-level log probabilities. Additionally, comparing each image against every class description is computationally expensive. To address these limitations, we propose an iterative MCQA approach that both reduces computational cost and eliminates the need for model logits access.

3.2 Iterative MCQA

We propose an *iterative MCQA* framework for zero-shot fine-grained image classification using LVLms. Given an image I and a set of N possible class descriptions, the LVLm iteratively selects a small subset of m options (e.g., $m = 5$) to progressively refine its prediction.

At test time, the model operates without access to the ground-truth label. In the first iteration, it is presented with a random subset of m class descriptions and selects the best-matching option. In subsequent iterations, the model retains the previously chosen option and compares it against $m-1$ new class descriptions that have not yet been evaluated. This process continues until all class descriptions are considered. The final selection in the last iteration is taken as the model's prediction. For example, in a 20-class task with $m = 5$, the model is first shown the set $\{A, B, C, D, E\}$. Suppose it selects A ; the next round compares A with $\{F, G, H, I\}$. If it then chooses G , the subsequent round presents $\{G, J, K, L, M\}$, and so on. The model's prediction is derived from its final selected option after iterating through the full set of classes.

While our iterative MCQA approach is entirely label-free, we incorporate an *early stopping mechanism* during evaluation for computational efficiency. Specifically, if the ground-truth class appears in a subset and the model fails to select it, we terminate the process early, as the final outcome will be incorrect regardless of subsequent iterations. This early termination serves purely as an evaluation shortcut and does not influence the model's decision-making process. Algorithm 1 (Appendix A.3) formalizes the complete iterative MCQA procedure.

3.3 Attention Intervention

Recent studies have demonstrated that LVLms exhibit a bias toward language tokens over image tokens during response generation (Liu et al., 2025; Wang et al., 2024a; Zhang et al., 2024). While existing methods address this through token logit manipulation (Gong et al., 2024; Huo et al., 2025) or decoder attention intervention (Liu et al., 2025; Zhou et al., 2025), they often rely on visual contrastive decoding techniques (Leng et al., 2024) that require more computation. These approaches, while effective for general tasks, prove computationally inefficient for fine-grained classification using LVLms where multiple class descriptions must be evaluated against each image.

Therefore, we propose a simple yet effective attention intervention technique that enhances visual information flow while maintaining computational efficiency. Our approach is motivated by two key findings: (1) LVLMs process critical visual information in intermediate early layers (Zhang et al., 2024), and (2) intermediate later/deep layers demonstrate better object-attribute grounding (Wang et al., 2024a; Jiang et al., 2025). Given an LVLM with L layers, we compute the average attention weights from layers 3 to k and incorporate them into layers $k + 1$ to $L - 2$, excluding the last two layers to optimize visual information flow. Additionally, we average the attention weights from layers $k + 1$ to $L - 2$ and propagate this grounding information to the final two layers:

$$\bar{A}_{\text{early}} = \frac{1}{k-2} \sum_{i=3}^k A_i \quad (3)$$

$$A'_j = A_j + \lambda \bar{A}_{\text{early}}, \quad j \in [k+1, L-2] \quad (4)$$

where A_i represents the attention weights in layer i , \bar{A}_{early} is the mean attention pattern from selected early layers, A'_j is the modified attention weights for layer j , and λ is a scaling factor controlling the influence of early visual information. With this approach, we attempt to ensure enhanced visual information flow and grounded object information while maintaining computational efficiency.

3.4 All-at-Once Approach

Recent advances in LVLM architectures have enabled processing of significantly longer contexts, with models like Qwen2.5-VL (Team, 2025), InternVL2.5 (Chen et al., 2024), and GPT-4o (OpenAI, 2024) supporting extended text prompts that can accommodate hundreds or thousands of tokens. This capability presents an alternative approach to our iterative MCQA framework by allowing all class descriptions to be presented simultaneously in a single forward pass.

In this approach, we provide the complete set of N class descriptions along with the input image and prompt the model to select the most appropriate description directly. The model receives a comprehensive view of all possible classes at once, potentially enabling more sophisticated reasoning about inter-class relationships and discriminative features. This method eliminates the iterative nature of our MCQA approach while maintaining the core principle of leveraging detailed class descriptions rather than generating class names.

We formulate this as a single-pass selection task where, given an image I and the complete set of N class descriptions $\{D_1, D_2, \dots, D_N\}$, the model is prompted to identify the index of the best-matching class description:

$$\hat{y} = \text{LVLM}(I, \{D_1, \dots, D_N\}, Q_{\text{all}}) \quad (5)$$

where Q_{all} represents a prompt that instructs the model to select the best matching class description from all provided options simultaneously.

4 Experiments and Evaluation

To evaluate our proposed method, we conduct comprehensive experiments across five fine-grained image classification benchmark datasets using different LVLMs and our curated class descriptions. This section presents the details of our class description curation process, experimental setup, evaluation results, and ablation studies.

4.1 Datasets

For our experiments, we utilize several established fine-grained image classification datasets: CUB (Wah et al., 2011) (200 bird species), Stanford Cars (Krause et al., 2013) (196 car models), FGVC Aircraft (Maji et al., 2013) (102 aircraft variants), Food-101 (Bossard et al., 2014) (101 food categories), and Stanford Dogs (Khosla et al., 2011) (120 dog breeds). These datasets present diverse challenges in fine-grained visual recognition across different domains including birds, dogs, cars, aircrafts, and foods. Following Kim and Ji (2024); Halbe et al. (2024), we conduct our experiments on randomly selected subsets of these datasets to maintain consistency with prior work, ensuring we exclude any images used in our class description curation process. For example, for the CUB dataset, we randomly select 5 images per class, resulting in 1,000 test images across 200 classes.

4.2 Class Description Curation

Existing class descriptions for these five datasets, provided by Menon and Vondrick (2022), Saha et al. (2024), and Kim and Ji (2024), are either generated using LLMs like GPT-3/4 (Menon and Vondrick, 2022; Saha et al., 2024) or extracted from the internet and refined using GPT-4V (Kim and Ji, 2024). We observe that these descriptions often lack precise fine-grained visual attributes necessary for accurate identification and may contain irrelevant information.




Images	Existing Class Descriptions	Ours Class Descriptions
 <p>Sooty Albatross</p>	<ul style="list-style-type: none"> Sooty-brown to sooty-black plumage. Darker head and tail. White ring around its eye. Yellow to orange line on its lower jaw. Glossy Black beak. 	<ul style="list-style-type: none"> Slender body. Dark Gray or Black Head Sharp black beak White eye-ring. Grayish-brown back and wings with darker tips. Long, narrow wings, slightly curved.
 <p>Crested Auklet</p>	<ul style="list-style-type: none"> Dark sooty grey body, wings, and tail. Reddish-orange bill with yellow tips. White auricular plumes. Grey legs and feet. Black claws. Black forward-curving crest feathers. 	<ul style="list-style-type: none"> Bright orange beak with a white stripe near the base. Distinctive tuft of feathers on the head. White eye with a black pupil. Dark gray to black body. Webbed feet, dark in color.
 <p>Red-winged Blackbird</p>	<ul style="list-style-type: none"> Black bird. Red shoulder patches. Pale yellow wing bar. 	<ul style="list-style-type: none"> Predominantly black plumage (head, beak, eyes, body). Distinctive bright orange and yellow patch on wings (near shoulder). Slender, conical beak.

Figure 3: Comparison of class descriptions from an existing dataset and our introduced class descriptions. Bold text in the "Ours Class Descriptions" column highlights key discriminative features that are either absent or described with less specificity in the "Existing Class Descriptions" column. The increased detail in our proposed descriptions facilitates more accurate zero-shot fine-grained image classification.

To address these limitations, we develop a two-stage process for curating more accurate class descriptions for these five datasets (prompts in Appendix A.2). First, we iteratively feed individual representative images from each class (10 images total) to Qwen2.5-VL-72B-Instruct (Team, 2025), a state-of-the-art LVM, to generate detailed visual descriptions for each image. Then, we process these 10 descriptions through Qwen2.5-72B-Instruct (Yang et al., 2024) LLM to compile a single comprehensive class description that captures common, distinctive visual features.

Our curation process yields 10 benchmarks - paired versions of 5 datasets with and without class names in their descriptions. Figure 3 demonstrates the enhanced descriptive precision of our curated descriptions for the CUB dataset (Wah et al., 2011) compared to existing class descriptions from Finer (Kim and Ji, 2024). Our descriptions incorporate more specific visual attributes (highlighted in bold) that are essential for fine-grained classification. For example, in describing the Crested Auklet, our description specifies a "bright orange beak with a white stripe near the base" and a "distinctive tuft of feathers", providing more precise detail than

the existing description's "reddish-orange bill with yellow tips" and "black forward-curving crest feathers". Similarly, for the Red-winged Blackbird, we include the "distinctive bright orange and yellow patch on wings (near shoulder)", offering crucial color and spatial information missing from the existing description's simple "red shoulder patches". This enhanced specificity across multiple visual features (beak morphology, plumage patterns, body structure) facilitates more accurate zero-shot classification among visually similar species.

4.3 Experimental Setup

For our experiments, we employ five state-of-the-art LVMs: LLaVA-v1.5-7B, LLaVA-v1.5-13B, InternVL2_5-8B, Qwen2-VL-7B-Instruct (also referred to as Qwen2-VL-7B in this literature), and Qwen2.5-VL-7B-Instruct. For our attention intervention experiments, we set $\lambda=1$ and $k=21$ for LLaVA-v1.5-7B, InternVL2_5-8B models and $k=23$ for LLaVA-v1.5-13B model. We evaluate model performance using classification accuracy as our primary metric.

4.4 Results

We evaluate our proposed approaches: Yes/No QA, iterative MCQA, iterative MCQA with attention intervention, and all-at-once across multiple LVMs on multiple fine-grained classification benchmarks. Table 1 presents comprehensive results under two experimental settings: with and without class names in class descriptions.

Comparative analysis of proposed methods. Our iterative MCQA approach consistently outperforms the Yes/No QA baseline across all models and experimental settings. In the "without name" setting, Qwen2.5-VL-7B-Instruct achieves 30.60% on CUB with iterative MCQA compared to 21.30% with Yes/No QA, representing a substantial 9.30% improvement. We observe a similar trend for Stanford Cars (41.63% vs 22.04%) and FGVC Aircraft (17.71% vs 10.57%) datasets with the Qwen2.5-VL-7B-Instruct model. The inclusion of class names dramatically shifts the performance landscape. Yes/No QA becomes highly competitive, with Qwen2-VL-7B-Instruct achieving exceptional results (53.30% on CUB, 86.84% on Stanford Cars, the latter being the best performance across all methods). Meanwhile, iterative MCQA maintains strong performance while offering the advantage of structured multi-step reasoning, particularly evident in models like InternVL2_5-8B achieving

Setting	Method	Model	CUB	Stanford Dogs	Stanford Cars	FGVC Aircraft	Food 101
W/o name	Yes/No QA	LLaVA-v1.5-7B	11.10	9.50	14.18	7.43	52.28
		LLaVA-v1.5-13B	13.80	10.50	18.67	9.71	56.24
		Qwen2-VL-7B	20.20	22.33	21.94	13.71	67.33
		Qwen2.5-VL-7B	21.30	18.00	22.04	10.57	65.15
	Iterative MCQA	LLaVA-v1.5-7B	9.00	15.67	9.90	11.71	40.99
		LLaVA-v1.5-7B + Attn	12.10	10.17	14.29	15.14	43.76
		Δ	+3.10	-5.50	+4.39	+3.43	+2.77
		LLaVA-v1.5-13B	7.50	16.00	13.88	10.57	32.87
		LLaVA-v1.5-13B + Attn	8.80	16.33	15.10	12.29	30.30
		Δ	+1.30	+0.33	+1.22	+1.72	-2.57
		InternVL2_5-8B	26.50	21.00	27.35	12.86	58.42
		InternVL2_5-8B + Attn	27.30	21.67	27.65	13.71	58.81
		Δ	+0.80	+0.67	+0.30	+0.85	+0.39
		Qwen2-VL-7B	32.40	15.67	33.78	20.00	54.85
		Qwen2.5-VL-7B	30.60	21.50	41.63	17.71	63.76
		All-at-Once	InternVL2_5-8B	15.30	17.67	16.22	10.57
Qwen2.5-VL-7B	6.60		14.33	21.94	12.00	43.76	
Qwen2.5-VL-32B	15.60		23.83	47.96	21.43	49.31	
GPT-4o	35.00		50.00	50.51	27.14	80.20	
o4-mini (reasoning)	48.50		58.33	64.29	40.00	89.11	
With name	Yes/No QA	LLaVA-v1.5-7B	15.60	17.00	36.22	19.14	55.84
		LLaVA-v1.5-13B	15.60	22.17	36.02	23.14	61.98
		Qwen2-VL-7B	53.30	62.33	86.84	63.43	75.45
		Qwen2.5-VL-7B	53.80	58.00	75.71	49.14	67.92
	Iterative MCQA	LLaVA-v1.5-7B	18.40	19.83	20.20	30.86	47.92
		LLaVA-v1.5-7B + Attn	23.30	22.50	21.94	31.14	48.32
		Δ	+4.90	+2.67	+1.74	+0.28	+0.40
		LLaVA-v1.5-13B	12.60	29.83	32.45	18.00	40.59
		LLaVA-v1.5-13B + Attn	13.80	31.33	31.43	18.57	37.43
		Δ	+1.20	+1.50	-1.02	+0.57	-3.16
		InternVL2_5-8B	36.60	38.50	42.65	33.43	56.63
		InternVL2_5-8B + Attn	37.50	38.00	42.86	32.57	57.03
		Δ	+0.90	-0.50	+0.21	-0.86	+0.40
		Qwen2-VL-7B	53.20	46.33	71.12	54.57	54.26
		Qwen2.5-VL-7B	53.00	46.17	65.92	52.86	67.13
		All-at-Once	InternVL2_5-8B	20.60	30.00	21.84	24.86
Qwen2.5-VL-7B	15.20		31.33	31.63	40.00	51.49	
Qwen2.5-VL-32B	33.00		52.67	62.24	42.86	55.05	
GPT-4o	70.50		74.17	80.10	77.14	90.10	
o4-mini (reasoning)	67.20		71.67	83.16	77.14	91.09	

Table 1: Zero-shot fine-grained visual classification performance in accuracy using different methods and models. Results compare Yes/No QA, iterative MCQA, and all-at-once approaches, with our proposed attention intervention (+Attn) enhancement. We conducted experiments with and without class names in descriptions on five fine-grained classification benchmarks. Δ rows show the performance gains from our introduced attention intervention. Best performance in each setting is highlighted in **bold**.

balanced performance across all datasets.

Impact of Attention Intervention. The attention intervention shows variable effectiveness, with improvements varying across architectures and datasets. LLaVA-v1.5-7B demonstrates the most consistent gains from attention intervention, particularly in the "with name" setting where it achieves +4.90% improvement on CUB (18.40% \rightarrow 23.30%). LLaVA-v1.5-13B shows moderate improvements ranging from +1.20% to +1.72% across most datasets, though with some minor decreases on Stanford Cars and Food-101. InternVL2_5-8B

exhibits the most stable response to attention intervention, with modest but consistent gains (+0.21% to +0.90%) across datasets in both settings. We performed layer selection optimization for the LLaVA-v1.5-7B model using the CUB dataset and applied a similar layer selection to the larger models. As we optimized layer selection specifically for LLaVA-v1.5-7B, the intervention proves most beneficial for this model. However, the fact that similar layer selection also shows performance gains for other models indicates that model-specific and dataset-specific layer selection could further enhance the

effectiveness of our attention intervention method.

All-at-once vs. iterative approach. Our all-at-once approach reveals interesting trade-offs between context utilization and performance. While state-of-the-art proprietary models (GPT-4o: 70.50% on CUB, o4-mini: 67.20% on CUB) achieve competitive results, open-source models suffer significant performance degradation when processing all class descriptions simultaneously. The stark contrast is evident in Qwen2.5-VL-7B-Instruct, which achieves 53.00% on CUB with iterative MCQA but only 15.20% with all-at-once in the "with name" setting, a 37.80% performance drop that highlights the cognitive burden of simultaneous multi-class comparison.

Key Findings. Four critical insights emerge from our evaluation: (1) Iterative MCQA consistently outperforms Yes/No QA, with average improvements ranging from 5% to 20% across datasets; (2) Attention intervention effectiveness varies across different model architectures, with LLaVA-v1.5-7B benefiting more than InternVL2_5-8B, particularly when using our optimized layer selection tailored for the LLaVA-v1.5-7B model on the CUB dataset; (3) Including class names significantly enhances performance across all approaches, with particularly notable improvements on FGVC Aircraft and Stanford Cars datasets; (4) The all-at-once approach works well for proprietary models but shows substantial degradation in open-source models, highlighting the advantage of iterative processing. These results demonstrate that structured, iterative processing with attention guidance provides superior performance for zero-shot fine-grained classification compared to single-pass approaches.

4.5 Comparison

Table 2 presents a performance comparison between our proposed method and existing approaches across four fine-grained visual classification datasets. Our method significantly outperforms Finer (Kim and Ji, 2024), the current SOTA approach, across all datasets. While Finer achieves accuracies of 1.56%, 4.20%, 0.50%, and 1.50% on CUB, Stanford Dogs, Stanford Cars, and FGVC Aircraft, respectively, our approaches demonstrate substantial improvements across all datasets.

Our iterative MCQA approach with attention intervention (MCQA + Attn) achieves SOTA performance on most datasets: CUB (23.30%), Stanford Dogs (22.50%) and Aircraft (31.14%), with

absolute accuracy gains of 21.74%, 18.30%, and 29.64% over Finer, respectively. Meanwhile, the Yes/No QA approach excels on Stanford Cars with the highest accuracy of 36.22%, achieving a substantial 35.72% improvement over the baseline.

These significant improvements, with absolute accuracy gains ranging from 18.30% to 35.72% over Finer, validate that our proposed approaches effectively enhance the model’s ability to focus on discriminative visual features for fine-grained classification. Among our proposed approaches, the iterative MCQA with attention intervention demonstrates the most consistent performance across multiple datasets. A detailed comparison with CLIP-based methods is provided in Appendix A.7, and a discussion of the computational efficiency of our approach relative to the baselines is included in Appendix A.6.

Method	CUB	Stanford Dogs	Stanford Cars	FGVC Aircraft
Finer (Kim and Ji, 2024)	1.56	4.20	0.50	1.50
Ours (Yes/No QA)	15.60	17.00	36.22	19.14
Ours (MCQA)	18.40	19.83	20.20	30.86
Ours (MCQA + Attn)	23.30	22.50	21.94	31.14

Table 2: Performance comparison of our proposed method with existing methods. For a fair comparison, all experiments use the LLaVA-v1.5-7B model. For Finer, values are extracted from the published results graphs with a potential variance of 1-2%. Best performance for each dataset is highlighted in **bold**.

4.6 Ablation Study

We conduct a series of ablation studies to evaluate the impact of key components in our proposed method. Additional ablation results are provided in Appendix A.4, and A.5.

4.6.1 Comparison of Datasets

Table 3 highlights the importance of our introduced class descriptions in zero-shot fine-grained classification using Qwen2.5-VL-7B-Instruct (Team, 2025) in the iterative MCQA setting. The existing class descriptions, sourced from the Finer (Kim and Ji, 2024) (except for Food-101, which comes from Menon and Vondrick (2022)), lack detailed attribute representation, leading to lower accuracy. In contrast, our introduced descriptions, which emphasize precise visual attributes, significantly improve performance, with notable gains on Food-101 (+13.07%) and CUB (+8.3%). This demonstrates that richer attribute-based descriptions enhance the model’s ability to distinguish fine-grained classes.

Including class names further improves accuracy across all datasets, as the model benefits from explicit textual cues. However, even in this setting, our descriptions consistently outperform the existing ones by up to +30.30%, demonstrating the effectiveness of structured visual attributes. Interestingly, the performance gap narrows with class names, suggesting that while names aid classification, detailed descriptions remain crucial for capturing fine-grained distinctions. This underscores the importance of precise textual descriptions for improving zero-shot classification in LVLMs.

Dataset	Without Class Name			With Class Name		
	Existing	Ours	Δ	Existing	Ours	Δ
CUB	22.30	30.60	+8.30	50.10	53.00	+2.90
Stanford Dogs	14.50	21.50	+7.00	44.00	46.17	+2.17
Stanford Cars	42.55	41.63	-0.92	65.20	65.92	+0.72
FGVC Aircrafts	14.86	17.71	+2.85	51.71	52.86	+1.15
Food-101	50.69	63.76	+13.07	36.83	67.13	+30.3

Table 3: Performance comparison of Qwen2.5-VL-7B-Instruct model with iterative MCQA approach on existing class descriptions and our introduced class descriptions, with and without the class name in the class descriptions. The Δ columns show the performance gain of our introduced class descriptions over the existing ones in both settings.

4.7 Impact of Model Scale

To evaluate the effect of model capacity, we conducted experiments using Qwen2.5-VL-7B-Instruct and Qwen2.5-VL-32B-Instruct in our all-at-once approach. As shown in Table 1, performance consistently improves with increased model size across most datasets and settings. In the "with name" setting, the 32B model demonstrates substantial improvements over the 7B variant: CUB (33.00% vs 15.20%), Stanford Dogs (52.67% vs 31.33%), Stanford Cars (62.24% vs 31.63%), FGVC Aircraft (42.86% vs 40.00%), and Food-101 (55.05% vs 51.49%). In the "without name" setting, similar trends are observed with significant improvements on Stanford Dogs (23.83% vs 14.33%) and Stanford Cars (47.96% vs 21.94%), while FGVC Aircraft shows notable enhancement (21.43% vs 12.00%). The results highlight that larger models generally provide better capacity for distinguishing fine-grained visual features, particularly when class names are available to provide additional semantic context. This scaling benefit is most pronounced on datasets requiring detailed visual discrimination, such as Stanford Cars and Stanford Dogs.

4.7.1 Robustness of Iterative MCQA

To assess the robustness of our iterative MCQA approach, we analyze how often models deviate from providing the requested index numbers by giving non-numerical responses. Our experiments show high instruction adherence across models. Among many experiments, LLaVA models produced only 2-43 non-numerical responses across 8 experiments, each involving 3,000-20,000 generation calls. Similarly, InternVL2_5-8B and Qwen2.5-VL models showed strong reliability with just 3-19 such errors in 3 experiments. These minimal error rates are negligible compared to the total number of model calls, and the rare instances of non-compliance could be addressed by directly computing probabilities for option indices. In contrast, the all-at-once approach showed more formatting issues, with Qwen2.5-VL and InternVL2_5-8B models occasionally providing non-numerical values or responses with explanatory text. The consistently low error rates in iterative MCQA validate that our approach is both efficient and robust.

5 Conclusion

The application of Large Vision-Language Models (LVLMs) to zero-shot fine-grained image classification represents a promising direction in vision-language understanding. We address current limitations by proposing multiple approaches: Yes/No QA, iterative multiple-choice question answering (MCQA), and all-at-once methods, enhanced by a simple yet effective attention intervention technique. Our investigation revealed the limitations of existing LLM-generated class descriptions, prompting us to curate refined descriptions that better capture fine-grained visual characteristics. Through comprehensive experiments across five datasets and multiple model architectures, we demonstrate that our iterative MCQA method consistently outperforms the Yes/No QA approach, while the all-at-once approach shows promising results for proprietary models but limitations for open-source architectures. Our methods achieve substantial performance improvements, significantly outperforming existing SOTA approaches across all evaluated datasets. The effectiveness of our attention intervention mechanism, particularly its benefits for smaller models, suggests promising directions in developing more sophisticated attention guidance techniques tailored to specific LVLm architectures for zero-shot fine-grained classification tasks.

6 Limitations

While our iterative MCQA approach with attention intervention demonstrates promising results in zero-shot fine-grained visual classification, several key limitations exist. The lower performance on most datasets indicates that current LVLMs still struggle with fine-grained visual distinctions. Additionally, our attention intervention technique’s effectiveness varies across architectures and datasets, suggesting dependency on model architecture and task characteristics. The quality of class descriptions is crucial, and despite leveraging advanced LVLMs and LLMs, they may still miss discriminative features that domain experts would identify. Refining these descriptions requires substantial expertise and manual effort, limiting scalability. Like other zero-shot approaches, our method’s effectiveness is bounded by the pre-training data quality and coverage of LVLMs, particularly for specialized domains underrepresented in the training data.

7 Acknowledgments

This project was supported in part by the Virginia Tech College of Science Academy of Data Science Discovery Fund (ADSDF) (Award: 238695) and the Research Acceleration Program (RAP) grant at Carilion Clinic, Roanoke, Virginia, USA. We also acknowledge Advanced Research Computing at Virginia Tech for providing computational resources and technical support that have contributed to the results reported within this paper. We thank all reviewers for their comments, which helped improve the paper.

References

- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, pages 446–461. Springer.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. 2025. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Mia Chiquier, Utkarsh Mall, and Carl Vondrick. 2024. Evolving interpretable visual classifiers with large language models. In *European Conference on Computer Vision*, pages 183–201. Springer.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. 2024. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv preprint arXiv:2409.17146*.
- Xuan Gong, Tianshi Ming, Xinpeng Wang, and Zhihua Wei. 2024. Damro: Dive into the attention mechanism of lvm to reduce object hallucination. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7696–7712.
- Shaunak Halbe, Junjiao Tian, KJ Joseph, James Seale Smith, Katherine Stevo, Vineeth N Balasubramanian, and Zsolt Kira. 2024. Grounding descriptions in images informs zero-shot visual recognition. *arXiv preprint arXiv:2412.04429*.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming Zhang, and Nenghai Yu. 2024. Opera: Alleviating hallucination in multi-modal large language models via over-trust penalty and retrospection-allocation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13418–13427.
- Fushuo Huo, Wenchao Xu, Zhong Zhang, Haozhao Wang, Zhicheng Chen, and Peilin Zhao. 2025. Self-reflective decoding: Alleviating hallucinations for large vision-language models. In *The Thirteenth International Conference on Learning Representations*.
- Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR.
- Nicholas Jiang, Anish Kachinthaya, Suzanne Petryk, and Yossi Gandelsman. 2025. Interpreting and editing vision-language representations to mitigate hallucinations. In *The Thirteenth International Conference on Learning Representations*.
- Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. 2011. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proc. CVPR workshop on fine-grained visual categorization (FGVC)*, volume 2.

- Jeonghwan Kim and Heng Ji. 2024. Finer: Investigating and enhancing fine-grained visual concept recognition in large vision language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6187–6207.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. 2020. Concept bottleneck models. In *International conference on machine learning*, pages 5338–5348. PMLR.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong Bing. 2024. Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13872–13882.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. 2024. Evaluating text-to-visual generation with image-to-text generation. In *European Conference on Computer Vision*, pages 366–384. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Shi Liu, Kecheng Zheng, and Wei Chen. 2025. Paying more attention to image: A training-free method for alleviating hallucination in vlms. In *European Conference on Computer Vision*, pages 125–140. Springer.
- Subhansu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*.
- Sachit Menon and Carl Vondrick. 2022. Visual classification via description from large language models. *arXiv preprint arXiv:2210.07183*.
- Michael Ogezi, Bradley Hauer, and Grzegorz Kondrak. 2024. Semantically-prompted language models improve visual descriptions. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 4285–4302.
- OpenAI. 2024. *Gpt-4o system card*. *Preprint*, arXiv:2410.21276.
- Sarah Pratt, Ian Covert, Rosanne Liu, and Ali Farhadi. 2023. What does a platypus look like? generating customized prompts for zero-shot image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15691–15701.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Oindrila Saha, Grant Van Horn, and Subhansu Maji. 2024. Improved zero-shot classification by adapting vlms with text descriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17542–17552.
- Qwen Team. 2025. *Qwen2.5-vl*.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The caltech-ucsd birds-200-2011 dataset. *California Institute of Technology*.
- Chenxi Wang, Xiang Chen, Ningyu Zhang, Bozhong Tian, Haoming Xu, Shumin Deng, and Huajun Chen. 2024a. Mllm can see? dynamic correction decoding for hallucination mitigation. *arXiv preprint arXiv:2410.11779*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024b. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Canshi Wei. 2024. Enhancing fine-grained image classifications via cascaded vision language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1857–1871.
- An Yan, Yu Wang, Yiwu Zhong, Chengyu Dong, Zexue He, Yujie Lu, William Yang Wang, Jingbo Shang, and Julian McAuley. 2023. Learning concise and descriptive attributes for visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3090–3100.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, and others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. 2023. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19187–19197.

Xiaofeng Zhang, Chen Shen, Xiaosong Yuan, Shaotian Yan, Liang Xie, Wenxiao Wang, Chaochen Gu, Hao Tang, and Jieping Ye. 2024. From redundancy to relevance: Enhancing explainability in multimodal large language models. *arXiv preprint arXiv:2406.06579*.

Guanyu Zhou, Yibo Yan, Xin Zou, Kun Wang, Aiwei Liu, and Xuming Hu. 2025. Mitigating modality prior-induced hallucinations in multimodal large language models via deciphering attention causality. In *The Thirteenth International Conference on Learning Representations*.

A Appendix

This section contains the details of the following topics:

- Prompt for iterative MCQA (Appendix A.1)
- Prompts for our Class Descriptions Generation (Appendix A.2)
- Algorithm of our Iterative MCQA (Appendix A.3)
- Impact of Number of Options in Iterative MCQA Approach (Appendix A.4)
- Impact of Prompts (Appendix A.5)
- Efficiency of the Iterative MCQA Framework (Appendix A.6)
- Comparison with CLIP-based Methods (Appendix A.7)

A.1 Prompt for iterative MCQA

For our proposed method, we use a consistent prompt template across all experiments. Figure 4 shows the prompt used in our iterative MCQA approach for zero-shot fine-grained image classification.

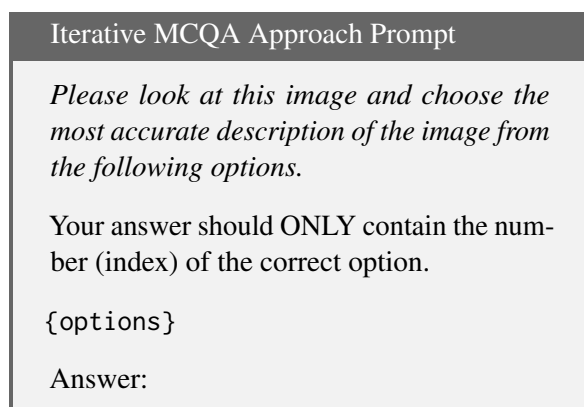


Figure 4: Prompt used for the iterative multiple-choice question answering (MCQA) approach.

Our prompt structure enforces methodological consistency across datasets and experimental settings while constraining model outputs to numerical indices.

A.2 Prompts for our Class Descriptions Generation

We developed a two-stage process for curating more accurate class descriptions:

First, we feed 10 representative images from each class to Qwen2.5-VL-72B-Instruct (Team, 2025), a state-of-the-art LVLM, to generate detailed visual descriptions for each image. Figure 5 shows used prompt for this step.

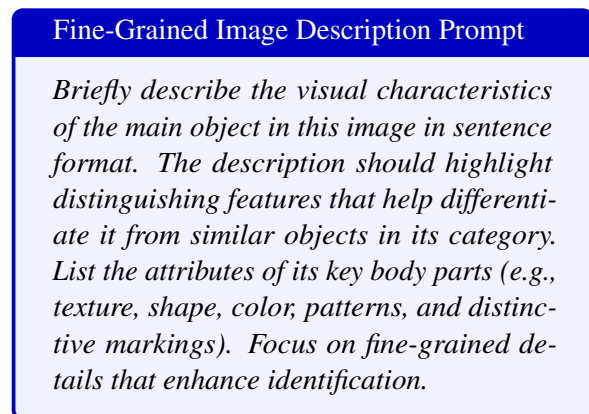


Figure 5: Prompt used for obtaining fine-grained visual descriptions from an LVLM.

Then, we process these 10 descriptions through Qwen2.5-72B-Instruct (Yang et al., 2024) LLM to synthesize a single comprehensive class description that captures common, distinctive visual features. Figure 6 depicts our employed prompt for this step. In the prompt, {image_descriptions} represents the 10 image descriptions we obtained from the previous step.

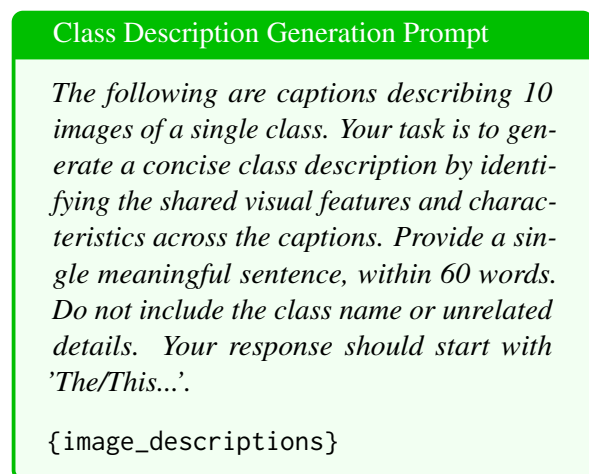


Figure 6: Prompt used for generating class descriptions from fine-grained image captions.

A.3 Algorithm of our Iterative MCQA

This section presents our introduced iterative MCQA algorithm. Algorithm 1 outlines the steps used in our zero-shot fine-grained classification task with LVLMs. Given an image I and a set of N possible class descriptions, the method iteratively selects a subset of m descriptions and queries the model to identify the best match. The process continues until either all class descriptions have been evaluated or a misclassification occurs, with the predicted class retained in the next $(m-1)$ descriptions. A prediction is considered correct only if the ground-truth class is consistently identified throughout all iterations (i.e., in the final prediction). Further details are provided in Section 3.2.

Algorithm 1 Iterative Multiple-Choice Question Answering for Zero-Shot Image Classification

Require: Image I , Set of Class Descriptions $C = \{c_1, c_2, \dots, c_n\}$, Model M , Batch Size m

Ensure: Predicted Class c^*

```
1: Initialize  $correct \leftarrow False$ ,  $processed \leftarrow False$ 
2: while not  $processed$  do
3:   Select  $m$  candidate descriptions  $C_m$ 
4:   Query  $M$  with  $I$  and  $C_m$ , obtaining predicted class  $c_p$ 
5:   if  $c \in C_m$  and  $c_p \neq c$  then
6:     return Misclassification  $\triangleright$  Terminate early as  $M$  missed the correct class
7:   else
8:     Move to next  $m$  descriptions keeping  $c_p$  in  $m$  descriptions.
9:   end if
10:  if all class descriptions are processed then
11:     $processed \leftarrow True$ 
12:  end if
13: end while
14: if  $c_p = c$  and  $processed$  then
15:  return Classified Correctly
16: else
17:  return Misclassified
18: end if
```

A.4 Impact of Number of Options in Iterative MCQA Approach

Traditional Multiple-Choice Question Answering (MCQA) approaches typically employ 4 or 5 options. Table 5 presents a comparative analysis of model performance when varying the number of multiple-choice options in our iterative MCQA

approach. We conducted experiments using the Qwen2.5-VL-7B-Instruct model across five fine-grained visual classification datasets. Our choice of using 5 options was motivated by three factors: alignment with traditional MCQA practices, the performance advantages observed with this configuration, and the context length limitations of the LLaVA-v1.5 models. As shown in the results, the impact of increasing options from 5 to 10 varies significantly across datasets. While some datasets show performance degradation (Stanford Dogs: -22.42%, FGVC Aircraft: -3.96%, Food-101: -8.12%), others demonstrate substantial improvements (Stanford Cars: +21.48%, CUB: +2.50%).

Notably, Stanford Cars shows the most dramatic improvement with 10 options (67.65% vs 46.17%), suggesting that for certain datasets with high inter-class similarity, additional options may help the model make more precise distinctions. Conversely, Stanford Dogs exhibits the largest performance drop (-22.42%), indicating that too many options can overwhelm the model’s decision-making process for certain classification tasks. With 5 options, the model achieved accuracies ranging from 46.17% to 67.13% across different datasets, while with 10 options, the range was 43.50% to 67.65%.

A.5 Impact of Prompts

We investigated the effect of prompt engineering on our iterative Multiple Choice Question Answering (MCQA) approach. Specifically, we enhanced the existing prompt by adding the instruction: *"The options are fine-grained. You have to utilize your grounding capability or visual attribute matching capability to select the best-aligned option."* This modification, designed to encourage more precise visual analysis, yielded mixed results when tested on our Qwen2.5-VL-7B-Instruct model.

For the Stanford Dogs dataset, we observed modest improvements: accuracy increased from 46.17% to 48.33% with class names provided, and from 21.50% to 22.67% without class names, respectively. A similar improvement was observed for the FGVC Aircraft dataset with class names, where accuracy improved from 52.86% to 54.57%. However, performance decreased for some datasets: the Food-101 dataset accuracy declined from 67.13% to 65.35% with class names.

An important trade-off emerged during testing: while the modified prompt generally improved accuracy, it also affected response consistency. Without the added instruction, the model consistently

Method	CUB	Food-101	Aircrafts	Cars
Ours	70.50	91.09	77.14	86.84
Menon and Vondrick (2022)	65.26	93.26	–	–
CuPL (Pratt et al., 2023)	–	93.33	36.69	76.49
CascadeVLM (Wei, 2024)	–	–	36.80	85.60
LaBo (Yang et al., 2023)	54.19	80.41	32.73	–
Yan et al. (Yan et al., 2023)	64.05	81.85	–	81.85
V-GLOSS (Ogezi et al., 2024)	–	–	38.70	–

Table 4: Performance comparison (Top-1 Accuracy %) of our method and prior CLIP-based methods across five fine-grained classification datasets. Missing values (–) indicate that the method did not report results for the corresponding dataset.

#Options	CUB	Stanford Cars	Stanford Dogs	FGVC Aircraft	Food 101
5	53.00	46.17	65.92	52.82	67.13
10	55.50	67.65	43.50	48.86	59.01

Table 5: Impact of the number of multiple-choice options on model accuracy across different fine-grained visual classification datasets. All experiments were conducted using the Qwen2.5-VL-7B-Instruct model with class descriptions including class names.

complied with the format requirement of returning only the option index. However, after introducing the new instruction, we observed that in 10 to 28 additional instances beyond our normal prompt, the model deviated from this format, providing additional explanatory text along with the index. This suggested that while the instruction helped improve accuracy in most cases, it compromised the model’s adherence to output formatting requirements.

A.6 Efficiency of the Iterative MCQA Framework

Our proposed iterative multiple-choice question answering (MCQA) framework provides a favorable balance between classification accuracy and computational efficiency. While traditional Yes/No QA approaches require $\mathcal{O}(N)$ forward passes for N class descriptions—one per class—our MCQA method reduces the number of passes to approximately $N/4$ by employing 5-way batching. For instance, in a dataset with 200 classes such as CUB, Yes/No QA requires 200 separate forward passes, whereas the iterative MCQA framework typically completes classification in only 40–50 passes, achieving a 75–80% reduction in computation. As shown in Table 5, increasing the batch size beyond 5 options results in decreased accuracy, indicating that 5-way MCQA provides a practical trade-off between inference cost and perfor-

mance. These results demonstrate that the iterative MCQA method enables scalable zero-shot fine-grained classification with significantly lower computational overhead.

A.7 Comparison with CLIP-based Methods.

We compare our zero-shot fine-grained classification framework with several recent CLIP-based and LVLm-augmented approaches across diverse datasets. Our results are based on the All-at-Once approach using GPT-4o or o4-mini. While traditional CLIP-based methods such as CuPL (Pratt et al., 2023) and Menon and Vondrick (2022) report high performance on Food-101 (93.33% and 93.26% respectively), they rely heavily on handcrafted prompts or multiple prompt augmentations. In contrast, our approach achieves competitive performance on Food-101 (91.09%) using a unified, training-free framework. On more fine-grained datasets, our method substantially outperforms LaBo (Yang et al., 2023) on CUB (70.50% vs. 54.19%) and FGVC Aircraft (77.14% vs. 32.73%), while achieving competitive results on Stanford Cars (83.16% vs. 81.85% from Yan et al. (2023)). Notably, CascadeVLM (Wei, 2024), which combines CLIP with LVLms, reports higher performance on Stanford Cars (85.60%) but much lower accuracy on FGVC Aircraft (36.80%), indicating limited generalizability. Our approach does not rely on any supervised classifier or handcrafted prompts, yet provides robust and consistent performance across multiple domains, highlighting its generalizability and effectiveness for zero-shot fine-grained image classification.